

Stream Reasoning-Based Control of Caching Strategies in CCN Routers^{*}

Harald Beck¹, Bruno Bierbaumer², Minh Dao-Tran¹, Thomas Eiter¹,
Hermann Hellwagner², and Konstantin Schekotihin²

¹ TU Wien, Vienna, Austria

{beck,dao,eiter}@kr.tuwien.ac.at

² Alpen-Adria-Universität Klagenfurt, Austria

bruno@itec.aau.at, firstname.lastname@aau.at

Abstract. Content-Centric Networking (CCN) research addresses the mismatch between the modern usage of the Internet and its outdated architecture. Importantly, CCN routers may locally cache frequently requested content in order to speed up delivery to end users. Thus, the issue of *caching strategies* arises, i.e., which content shall be stored and when it should be replaced. In this work, we employ novel techniques towards intelligent administration of CCN routers that autonomously switch between existing strategies in response to changing content request patterns. In particular, we present a router architecture for CCN networks that is controlled by rule-based stream reasoning, following the recent formal framework LARS which extends Answer Set Programming for streams. The obtained possibility for flexible router configuration at runtime allows for faster experimentation and may thus help to advance the further development of CCN. Moreover, the empirical evaluation of our feasibility study shows that the resulting caching agent may give significant performance gains.

1 Introduction

The architecture of the Internet is rooted in research on packet switching in the late 1960s and early 1970s. It evolved from a small research network that focused on sending text messages, to a global network of content distribution; Cisco estimates that by 2019, 80% of the Internet traffic will be video content [8]. However, the architectural foundation dates back to the 1980s, where massive data volumes and scalability were no concern. Commercial Content Distribution Networks (CDNs) have been developed as a workaround to cope with today's fast content delivery demands, which are built as overlays on the traditional Internet architecture (TCP/IP). In general, today's Internet architecture does not fit applications and uses well that have evolved meanwhile [17].

In response to this, various Future Internet research efforts are being pursued, among them Information-Centric Networking (ICN) [28], and in particular *Content-Centric Networking* (CCN) [18]. CCN attempts to replace the current location-based addressing

^{*} This work was partly funded by the Austrian Science Fund (FWF) under the CHIST-ERA project CONCERT (A Context-Adaptive Content Ecosystem Under Uncertainty), project number I1402

with a name/content-based approach. That is, data packets shall be routed and retrieved based on *what* the user wants, not from *where* it is retrieved. In a sense, CDN provides this, yet CCN supports this at the network level by making content identifiable.

An important element of the CCN architecture is that every CCN router has a *cache* (content store) which holds content items that were recently transmitted via the router. A request for a content item may be satisfied by a router rather than routed to the original content source; thus data is delivered to end users faster. A *caching strategy* defines which content is stored, on which routers, and for how long before being replaced. There is a rich literature of strategies and mechanisms for ICN/CCN [31,27,7,4], with a variety of parameters influencing the overall behavior of a network.

The caching strategies can be roughly classified into adaptive and reactive ones. The adaptive strategies use information about interests of users saved by a network logging system. This information is then used to estimate popularity of content in the future and push it to the caches of routers. Therefore, adaptive strategies are mostly used in CDNs which, by their nature, are tightly integrated with the networks of content providers. Strategies used in CCNs are essentially reactive, i.e. they use a kind of heuristic to predict whether a forwarded content chunk might be interesting for other users. If yes, the chunk is added to the router's cache. Some of the reactive strategies go even further and allow for synchronization of caching decisions between multiple routers. For instance, most popular content must be cached by routers of the lowest levels in the network topology. Such strategies, however, often work only for specific topologies, like trees.

Recent evaluations of CCN caching strategies, like [31], indicate that no “silver bullet” strategy is superior in all tested scenarios. Furthermore, selecting a good caching strategy and fine-tuning its parameters is difficult, as the distribution of consumer interests in content may vary greatly over time [29,6].

Example 1 Consider a situation in which some music clips go viral, i.e., get very popular over a short period of time. In this case, network administrators may manually configure the routers to cache highly popular content for some time period, and to switch back to the usual caching strategy when the consumer interests get more evenly distributed. However, as this period of time is hard to predict, it would be desirable that routers autonomously switch their caching strategy to ensure high quality of service. ■

As real CCNs are not deployed yet, there is currently no real-world experience to rely on, and developing selection methods for caching strategies is not well supported.

Motivated by all this, we consider a router architecture that allows for dynamic switching of caching strategies in reaction to the current network traffic, based on *stream reasoning*, i.e., reasoning over recent snapshots of data streams.

Contributions. Our contributions can be summarized as follows.

(1) We present an *Intelligent Caching Agent* (ICA) for the administration of CCN routers using stream reasoning, with the following features:

- ICA extends a typical CCN architecture with a decision unit, resulting in the first implementation of a local and dynamic selection of an appropriate caching strategy.
- The main component of the decision unit is based on the rule-based stream reasoning framework LARS [3], which is an extension of Answer Set Programming (ASP) for streams (see Section 3 for details). Implemented as a DLVHEX [12] plug-in, it

enables administrators to control caching strategy selection in a concise and purely declarative way. Furthermore, the selection control can be modified without taking a router offline, which is another important criterion for such systems.

(2) To support the development and testing of dynamic cache strategy selection, we propose an extension of ndnSIM [20] – a well-known CCN simulator – for iterative empirical assessment of proposed solutions for intelligent administration. In particular, the extension is designed to: (i) simulate various CCN application scenarios, (ii) implement different architectures of CCN routers, (iii) apply rule-based stream reasoning to make decisions about the caching strategy configuration for every router in the network, (iv) react quickly to inferred information from continuously streaming data and (v) be expressible in an understandable and flexible way for fast experimentation.

(3) We provide a detailed evaluation of our methods on two sample scenarios in which content consumers unexpectedly change their interests, as in Example 1. Our results indicate a clear performance gain when basic caching strategies are dynamically switched by routers in reaction to the observed stream of requested data packets.

In summary, we provide a feasibility study for using logic-based stream reasoning techniques to guide selection of caching strategies in CCNs. Moreover, we also provide a detailed showcase of analytical, declarative stream reasoning tools for intelligent administration problems; to the best of our knowledge, no similar work exists to date.

2 Content-Centric Networking

The operation of a CCN network relies on two packet types, *Interest* and *Data* packets. Clients issue *Interest* packets containing the *content name* they want to retrieve. CCN routers forward the *Interest* packets until they reach a content provider that can satisfy them with the content addressed by the *content name*. The content provider answers with a *Data* packet which travels back to the original content consumer following the previous *Interest* packets. In addition to delivering the *Data* packets back to the consumer, the CCN routers have the possibility to cache these packets in their *Content Stores*. Thus, the *Interest* packets of another consumer can be directly satisfied out of a *Content Store* without the need of going all the way to the original content provider. These caches make it possible to keep popular content near the consumer, satisfy content requests directly out of caches and reduce the network load [18].

Content Popularity Distribution. Not all content is equally popular. Usually, there is a small number of very popular content items and lots of unpopular ones, which is described in the literature with a *Zipf* distribution [23]. Let C be a number of items in the content catalog, α be a value of the exponent characterizing the distribution and i be a rank of an item in the catalog. Then, Zipf distribution predicts the frequency of Interest packets for item i as [24]:

$$P(X = i) = \frac{1/i^\alpha}{\sum_{j=1}^C 1/j^\alpha} \quad (1)$$

The variation of the exponent α allows to characterize different popularity models for contents requested by consumers: (i) if α is high, the popular content is limited to a small number of items; (ii) if α is low, every content is almost equally popular.

The content popularity distribution and its exponent α can be estimated by counting the *Interest* packets arriving at a router. The estimated values $\hat{\alpha}$ of the α parameter can be used to form rules like: “If a small number of content items has been very popular ($\hat{\alpha} \geq 1.8$) for the last 5 minutes, then action C should be applied.”

Caching strategies. A caching strategy decides which item gets replaced in the full cache storage if a new item should be added. We consider the following strategies [20]:

- *Least Recently Used.* The LRU strategy keeps the cached items in a list sorted by their access time stamps and replaces the oldest item.
- *First-In-First-Out.* For FIFO strategy, the cache is implemented as a simple queue and replaces the earliest inserted item.
- *Least Frequently Used.* The LFU strategy counts how often an item in the cache is accessed. When caching a new item, the item with the smallest access count is replaced.
- *Random.* The Random strategy replaces a random item in the cache with a new one.

3 Stream Reasoning

Stream reasoning [9] emerged from stream processing for real-time reasoning about information from data streams. Initially, the focus was on *continuous* queries akin to SQL [2,1]. Later works also dealt with advanced logic-oriented reasoning [15,14,30,21] on streaming data. In particular, LARS [3] has been proposed for stream-oriented logical reasoning in the spirit of Answer Set Programming (ASP) [5].

To the best of our knowledge, stream reasoning has not yet been considered in CCN as such. We argue that, from an information-oriented point of view, CCN is to a large degree a task of stream processing. In particular, the intelligent cache administration of routers adds the need to logically reason over the streaming data in real-time.

Example 2 (con’t) Consider the following rules to select a caching strategy. If in the last 30 seconds there was always a high $\hat{\alpha}$ value (some content is very popular), use LFU, and for a medium value, take LRU. Furthermore, use FIFO if the value is low but once in the last 20 seconds 50% was real-time content. Otherwise, use Random. ■

Example 2 illustrates that a fully declarative, rule-based language would assist the readability of a router’s module that controls such decisions. Moreover, it would allow administrators to update the control unit on-the-fly, i.e., without taking a router offline.

Notably, envisaged deployments of CCNs will involve more complex rules, where advanced reasoning features will be beneficial. This includes declarative exception handling, reasoning with multiple models, defaults, and the possibility to adjust the involved logic in a flexible, elaboration-tolerant and modular way. E.g., further information such as Data packet headers, network behavior and router internals can be accounted for in a comprehensible way by adding or adjusting rules using new predicates.

On top of such features, as offered by ASP, LARS provides operators to deal with stream-specific information, i.e., to access to temporal information and the possibility to limit reasoning to recent *windows* of data. Such recent snapshots of data can also be expressed in traditional stream processing languages like CQL [1]. While SQL-like languages provide a high degree of declarativity, more complex decision-making quickly becomes unreadable (because nested) and is less modular than rule-based approaches.

Moreover, administrators are usually familiar with rule-based configuration from other tools like IP tables. Therefore, stating decision-making processes as small if-then statements is more natural than encoding them in (often complex) SQL queries. Furthermore, updates at runtime allow for manual interventions in novel network situations. The essence of new situations, as understood by human administrators, can be added to the existing knowledge base without the need for recompiling the entire system.

While we do not elaborate here on multiple models or preferences, exploiting further advantages of ASP is suggestive but remains for subsequent work. For instance, as our implementation is based on DLVHEX (see Section 4), enriching the reasoning system e.g. with access to ontologies for object classification or to a scheduler is easy.

In this work, we focus on simple examples to convey core ideas and illustrate some of the benefits of stream reasoning within a specific simulation architecture. However, we emphasize that LARS as such provides a high degree of expressivity and is suitable for more involved setups which may build on this feasibility study.

3.1 Streams and Windows

Central to reasoning with LARS is the notion of a stream, which associates atoms with time points. Throughout, we distinguish *extensional atoms* \mathcal{A}^E for input data and *intensional atoms* \mathcal{A}^I for derived information. By $\mathcal{A} = \mathcal{A}^E \cup \mathcal{A}^I$ we denote the set of *atoms*.

Definition 1 (Stream) A stream $S = (T, v)$ consists of a timeline T , which is a closed interval $T \subseteq \mathbb{N}$ of integers called time points, and an evaluation function $v: \mathbb{N} \mapsto 2^{\mathcal{A}}$.

We call $S = (T, v)$ a *data stream*, if it contains only extensional atoms. We say that the timeline ranges from t_1 to t_2 , if $T = [t_1, t_2]$. To cope with the amount of data, one usually considers only recent atoms. Let $S = (T, v)$ and $S' = (T', v')$ be two streams s.t. $S' \subseteq S$, i.e., $T' \subseteq T$ and $v'(t') \subseteq v(t')$ for all $t' \in T'$. Then S' is called a *substream* or *window* of S . We may *restrict* an evaluation function v to a timeline T , defined as $v|_T(t) = v(t)$, if $t \in T$, else \emptyset .

By a *window function* w we understand a function that takes as input a stream $S = (T, v)$, a time point $t \in T$, and returns a window $S' \subseteq S$. Typical are *tuple-based* window functions that collect the most recent atoms of a given number, and *time-based* window functions that select all atoms of a given temporal range. In general, time-based windows move along the timeline in steps of a given size $d \geq 1$. For instance, if $d = k$, where k is the length of the window, one gets a *tumbling* time-based window function. In this work, we only use *sliding* time-based window functions $\tau(k)$ which always return the window of the most recent k time points, i.e., $d = 1$.

Definition 2 (Sliding Time-based Window) Let $S = (T, v)$ be a stream, $t \in T = [t_1, t_2]$ and $k \in \mathbb{N}$. Moreover, let $T' = [t', t]$ such that $t' = \max\{t_1, t - k\}$. Then, the (sliding) time-based window (of size k) is defined by $\tau(k)(S, t) = (T', v|_{T'})$.

Example 3 Consider a stream with a timeline $T = [0, 1800]$ that contains two atoms, indicating that, at time 42 and 987, at least 50% of all Data packets were real-time content. This is formalized by $S = (T, v)$, where $v(42) = v(987) = \{rtm50\}$, and $v(t) = \emptyset$ for all $t \in T \setminus \{42, 987\}$. The time-based window of size 30 at $t = 70$ is

defined as $\tau(30)(S, 70) = ([40, 70], v')$, where $v'(42) = \{rtm50\}$ and $v'(t') = \emptyset$ for all $t' \in [40, 70] \setminus \{42\}$. ■

3.2 LARS Formulas

Syntax. LARS adds new operators to propositional formulas.

- *Window operators* \boxplus^w . Formula evaluation in LARS is always relative to a time point $t \in T$ in the scope of the currently considered window $S = (T, v)$ (initially the entire stream). For every window function w , employing an expression $\boxplus^w \varphi$ will restrict the evaluation of the formula φ to the window obtained by $w(S, t)$.
- *Temporal quantification with* \Diamond and \Box . Often, one is interested whether a formula φ holds at *some* time point in a selected window, or at *all* time points. This is expressed by $\Diamond \varphi$ and $\Box \varphi$, resp.
- *Temporal specification with* $@_{t'}$. Dually, the $@$ operator allows to ‘jump’ to a specific time point t' (within T). That is to say, $@_{t'} \varphi$ evaluates φ at time point t' .

Based on these ingredients for dealing with information that is specific for streams, we define LARS formulas as follows.

Definition 3 (Formulas) Let $a \in \mathcal{A}$ be an atom and $t \in \mathbb{N}$. The set \mathcal{F} of formulas is defined by the grammar $\varphi ::= a \mid \neg \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \Diamond \varphi \mid \Box \varphi \mid @_t \varphi \mid \boxplus^w \varphi$.

Since we only use time-based window functions, we define that \boxplus^k abbreviates $\boxplus^{\tau(k)}$, i.e., \boxplus^k employs $\tau(k)$, a time-based window function of size k .

Example 4 The implication $\varphi = \boxplus^{30} \Box high \rightarrow use(lfu)$ informally says that if in the last 30 seconds (\boxplus^{30}) the predicate *high* always (\Box) holds, use *lfu*. ■

Semantics. In addition to streams, we consider background knowledge in form of a static data set, i.e., a set $\mathcal{B} \subseteq \mathcal{A}$ of atoms. From a semantic perspective, the difference to streams is that background data is always available, regardless of window applications.

Definition 4 (Structure) Let $S = (T, v)$ be a stream, W be a set of window functions and $\mathcal{B} \subseteq \mathcal{A}$ a set of facts. Then, we call $M = \langle S, W, \mathcal{B} \rangle$ a structure, S the interpretation stream and \mathcal{B} the background data of M .

Throughout, we will assume that W is the set of all time-based window functions.

Definition 5 (Entailment) Let $S^* = (T^*, v^*)$ be a stream, $S = (T, v)$ be a substream of S^* , and let $M = \langle S^*, W, \mathcal{B} \rangle$ be a structure. Moreover, let $t \in T$. The entailment relation \models between (M, S, t) and formulas is defined as follows. Let $a \in \mathcal{A}$ be an atom, and let $\varphi, \psi \in \mathcal{F}$ be formulas. Then,

$$\begin{array}{lll}
M, S, t \models a & \text{if } a \in v(t) \text{ or } a \in \mathcal{B}, & M, S, t \models \Diamond \varphi \quad \text{if } M, S, t' \models \varphi \text{ for some } t' \in T, \\
M, S, t \models \neg \varphi & \text{if } M, S, t \not\models \varphi, & M, S, t \models \Box \varphi \quad \text{if } M, S, t' \models \varphi \text{ for all } t' \in T, \\
M, S, t \models \varphi \wedge \psi & \text{if } M, S, t \models \varphi \text{ and } & M, S, t \models @_{t'} \varphi \text{ if } M, S, t' \models \varphi \text{ and } t' \in T, \\
& M, S, t \models \psi, & M, S, t \models \boxplus^w \varphi \text{ if } M, S', t \models \varphi \text{ where} \\
M, S, t \models \varphi \vee \psi & \text{if } M, S, t \models \varphi \text{ or } & S' = w(S, t). \\
& M, S, t \models \psi, & \\
M, S, t \models \varphi \rightarrow \psi & \text{if } M, S, t \not\models \varphi \text{ or } M, S, t \models \psi, &
\end{array}$$

$$\begin{aligned}
r_1 : @_T high &\leftarrow \boxplus^{30} @_T \hat{\alpha}(V), V \geq 18. & r_5 : use(lru) &\leftarrow \boxplus^{30} \square mid. \\
r_2 : @_T mid &\leftarrow \boxplus^{30} @_T \hat{\alpha}(V), 12 \leq V < 18. & r_6 : use(fifo) &\leftarrow \boxplus^{30} \square low, \boxplus^{20} \diamond rtm50. \\
r_3 : @_T low &\leftarrow \boxplus^{30} @_T \hat{\alpha}(V), V < 12. & r_7 : done &\leftarrow use(lfu) \vee use(lru) \vee use(fifo). \\
r_4 : use(lfu) &\leftarrow \boxplus^{30} \square high. & r_8 : use(random) &\leftarrow not\ done.
\end{aligned}$$

Fig. 1: Program P deciding which caching strategy to use

If $M, S, t \models \varphi$ holds, we say that (M, S, t) *entails* φ . Moreover, M *satisfies* φ at time t , if (M, S^*, t) entails φ . In this case we write $M, t \models \varphi$ and call M a *model* of φ at time t . Satisfaction and the notion of a model are extended to sets of formulas as usual.

Example 5 (cont'd) Consider $S = ([0, 1800], v)$, where *high* holds from sec 600 to 1200, i.e., $high \in v(t)$ for all $t \in [600, 1200]$. We evaluate φ from Example 4 at $t = 750$, i.e., the entailment $M, S, 750 \models \varphi$. The window operator \boxplus^{30} selects the substream $S' = (T', v|_{T'})$, where $T' = [720, 750]$. Clearly, $M, S', t' \models high$ for all $t' \in T'$ and thus $M, S', 750 \models \square high$. Hence, $M, S, 750 \models \varphi$ holds iff $M, S, 750 \models use(lfu)$ holds. ■

3.3 LARS Programs

LARS programs extend ASP [5], using the FLP-reduct [13], where rule literals can be replaced by LARS formulas.

Syntax. A *rule* r is an expression of form $\alpha \leftarrow \beta(r)$, where $H(r) = \alpha$ is the *head* and $\beta(r) = \beta_1, \dots, \beta_j, not\ \beta_{j+1}, \dots, not\ \beta_n, n \geq 0$, is the *body* of r . Here, $\alpha, \beta_1, \dots, \beta_n \in \mathcal{F}$ and all predicates in α are intensional. A (LARS) *program* P is a set of rules.

Example 6 Fig. 1 presents a formalization of the rules given in Example 2. Note that rule (r_4) corresponds to formula φ of Example 4. Rule (r_1) serves to derive *high* for each second T (in the selected interval) where value V in predicate $\hat{\alpha}$ is at least 1.8 (represented by integer 18). Thus, atom *high* abstracts away the specific value, and as long as it is always above the threshold of 1.8 during interval, atom $use(lfu)$ shall be concluded. Expressions like “ $V \geq 18$ ” are syntactic sugar for predefined predicates that are assumed to be included in the background data \mathcal{B} . ■

Note that variables used in Example 6 schematically abbreviate according ground rules. We give a formal semantics for the latter.

Semantics. For a data stream $D = (T_D, v_D)$, any stream $I = (T, v) \supseteq D$ that coincides with D on $\mathcal{A}^\mathcal{E}$ is an *interpretation stream* for D . Then, a structure $M = \langle I, W, \mathcal{B} \rangle$ is an *interpretation* for D . We assume W and \mathcal{B} are fixed and thus also omit them.

Satisfaction by M at $t \in T$ is as follows: $M, t \models \varphi$ for $\varphi \in \mathcal{F}$, if φ holds in I at time t ; $M, t \models r$ for rule r , if $M, t \models \beta(r)$ implies $M, t \models H(r)$, where $M, t \models \beta(r)$, if (i) $M, t \models \beta_i$ for all $i \in \{1, \dots, j\}$ and (ii) $M, t \not\models \beta_i$ for all $i \in \{j+1, \dots, n\}$; and $M, t \models P$ for program P , i.e., M is a *model* of P (for D) at t , if $M, t \models r$ for all $r \in P$. Moreover, M is *minimal*, if there is no model $M' = \langle I', W, \mathcal{B} \rangle \neq M$ of P s.t. $I' = (T, v')$ and $v' \subseteq v$. Note that smaller models must have the same timeline.

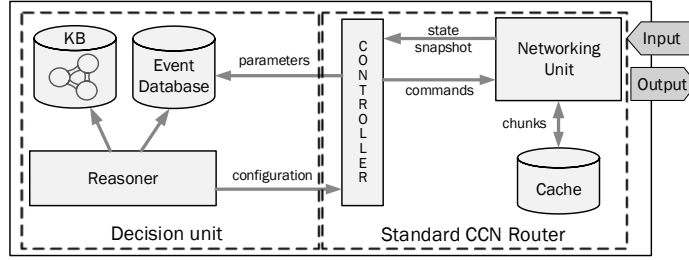


Fig. 2: Architecture of an Intelligent Caching Agent (ICA)

Definition 6 (Answer Stream) Let D be a data stream. An interpretation stream $I \supseteq D$ is an answer stream of program P for D at time t , if $M = \langle I, W, \mathcal{B} \rangle$ is a minimal model of the reduct $P^{M,t} = \{r \in P \mid M, t \models \beta(r)\}$. By $\mathcal{AS}(P, D, t)$ we denote the set of all such answer streams I .

Example 7 (cont'd) Consider a data stream $D = ([0, 1800], v_D)$, such that for every $t \in [600, 1200]$, there exists exactly one integer $V \geq 18$ s.t. $v_D(t) = \{\hat{\alpha}(V)\}$, and $v_D(t) = \emptyset$ for all $t \in [0, 599] \cup [1201, 1800]$. We evaluate program P of Fig. 1 at $t' = 750$. Clearly, the body of rule (r_1) holds at t' . Thus, to satisfy (r_1) , we need an interpretation stream $I = ([0, 1800], v)$ for D that also contains *high* in $v(t)$ for all $t \in [720, 750]$. Then $\boxplus^{30} \square \text{high}$ holds at t' , so *use(lfu)* must hold (at t') due to rule (r_4) . Now rule (r_7) requires *done* $\in v(t')$; which invalidates the body of (r_8) . If $I \supseteq D$ contains exactly these additions, it is the unique answer stream of P for D at t' . ■

4 System Description

As shown in Fig. 2, an Intelligent Caching Agent (ICA) extends the architecture of a common CCN router comprising a *networking unit* with a number of communication interfaces. This unit is responsible for the basic functionality of the router such as processing, forwarding of packets, etc. The networking unit is observed by a *controller*, which implements various supervising functions including a number of caching strategies. During operation of a router, the networking unit consults the selected caching strategy of the controller to identify video chunks to be stored in the cache. If the cache is full, the strategy also decides which cached chunk to replace. Given a non-empty cache, the networking unit for every arriving *Interest* packet checks whether it can be answered with cached chunks.

Decision unit. The decision unit of an ICA consists of three main components: (1) a database (DB) storing snapshots of parameters observed by the controller, (2) a knowledge base (KB) containing the ICA logic and (3) a reasoner that decides about configuration of the controller given the KB and a series of events in the DB.

The components (2) and (3) are based on the LARS framework, which we implemented using DLVHEX 2.5 [12]. The language of this system, Higher-order logic programs with EXternal atoms (HEX-programs) [11], is like LARS an extension of

```

1 intvl(S,E) :- &getSolverTime[](E), S=E-30.
2 intv2(S,E) :- &getSolverTime[](E), S=E-20.

3 val(high,S,E,T) :- &w[S,E,alpha](T,V), V>=18, intvl(S,E).
4 val(mid,S,E,T) :- &w[S,E,alpha](T,V), 12<=V, V<18, intvl(S,E).
5 val(low,S,E,T) :- &w[S,E,alpha](T,V), V<12, intvl(S,E).
6 val(rtm50,S,E,T) :- &w[S,E,rtc](T,V), V>50, intv2(S,E).

7 some(ID,S,E) :- val(ID,S,E,_).
8 always(ID,S,E) :- val(ID,S,E,_), val(ID,S,E,T):T=S..E.

9 use(lfu) :- always(high,S,E), intvl(S,E).
10 use(lru) :- always(mid,S,E), intvl(S,E).
11 use(fifo) :- always(low,S,E), intvl(S,E), some(rtm50,S2,E2), intv2(S2,E2).

12 done :- use(X), X!=random.
13 use(random) :- not done.

```

Listing 1.1: DLVHEX encoding for ICA

ASP [16] interpreted under FLP-semantics [13]. Formally, a HEX-program is a set of rules of the form $\alpha \leftarrow \beta(r)$, where α is a higher-order atom and $\beta_1, \dots, \beta_n \in \beta(r)$ are higher-order atoms or external atoms. A *higher-order atom* is a tuple $Y_0(Y_1, \dots, Y_n)$ where Y_0, \dots, Y_n are terms. An *external atom* has the form $\&g[Y_1, \dots, Y_n](X_1, \dots, X_m)$, where all X_i and Y_j are terms and $\&g$ is an external predicate name.

According to the semantics of HEX-programs, for every external predicate there is an external computation function such that the tuples (Y_1, \dots, Y_n) and (X_1, \dots, X_m) correspond to the input and output of the function, resp. Thus, HEX solvers allow for a bidirectional information flow between the solver and an external oracle. For instance, in the DLVHEX system external functions are defined as solver plug-ins which implement the semantics of external atoms. In our case external atoms can be used to get information about events stored in the database, compute required statistics, etc.

Our implementation defines an external atom $\&w[S, E, F](T, V)$ representing the described time-based LARS window operator. The terms $S, E \in \mathbb{N}$ define the time interval of the window function (Definition 2) and F is a string comprising a function name. Our DLVHEX plug-in evaluates the function over events registered in the database within the given time interval and returns its results as a set of tuples $\{(t_1, v_1), \dots, (t_k, v_k)\}$, where t_i and v_i indicate the time point and the value of a function, respectively. To define rules that respect only recent events, we use an external atom $\&getSolverTime[](E)$; it has no inputs and outputs the current system time E .

The DLVHEX encoding for ICA is presented in Listing 1.1, which corresponds to the LARS encoding presented in Fig. 1 and could be in principle automatically generated from it. Rules 1 and 2 derive time intervals for which the reasoning must be done. The next three rules find all time points in the last 30 seconds in which a router found that the content popularity is high, medium or low. That is, for the estimated value $\hat{\alpha}$ of the parameter α of the Zipf distribution (1) we have either $\hat{\alpha} \leq 1.8$, or $1.2 \leq \hat{\alpha} < 1.8$, or $\hat{\alpha} < 1.2$. The selection of value intervals was done empirically and its choice depends on the desired sensitivity of the ICA to changing conditions. Note that the parameter values may also be selected using machine learning techniques. Atoms of the form $val(ID, S, E, T)$ indicate that an event ID was registered at the point T of the time




Fig. 3: Simulation architecture of an Intelligent Caching Agent

interval $[S, E]$. Rule 6 derives all time points from the last 20 secs for which 50% of all *Interest* packets asked for real-time content, like broadcasts, video calls, etc.

The LARS \diamond and \square operators are represented by rules 7 and 8. The former is used to derive that the event ID occurred in some time point of the interval $[S, E]$, whereas the latter indicates that an event occurred at all time points of the interval. Note that the operator “:” in rule 8 generates a conjunction of atoms, where the operator “..” iteratively assigns every number $n \in [S, E]$ to the variable T .

Finally, the remaining rules implement the caching strategy selection. A router is configured to use one of the strategies – LFU, LRU or FIFO – when corresponding pre-conditions are fulfilled. Otherwise, the random caching strategy is selected by default.

Simulation environment. We implemented our ICA approach by extending the CCN simulator ndnSIM 2.0 [20] as shown in Fig. 3; a Content Store Tracer component was added to observe states of the router components of the simulator and push this data to the event database. Similarly to [10], our extension of ndnSIM periodically triggers the solving process for a decision about the controller configuration based on the events stored in the database. The process invokes the DLVHEX solver which takes `solver.hex`, the HEX-program in Listing 1.1, and the system time as input. In addition, the solver consults `solver.py`, a Python script that implements evaluation of external predicates, access to the events database, functions, like `alpha` or `rtc`.




Fig. 4: Abilene network topology [22]

5 Evaluation

Section 4 proposed an agent-based caching system, making use of a decision component written in LARS. We now present the evaluation of the resulting simulation system presented in Fig. 3. We show the applicability of our architecture for dynamic caching and demonstrate the potential performance gains over static caching approaches.

5.1 Setup

The evaluation setup consists of four main parts: the chosen network topology, the considered scenarios of user behaviour, the employed caching strategies and the description of the system parameters that influence performance.

Network. As mentioned before, much further research needs to be carried out before CCNs can be deployed in real-world applications. Consequently, empirical CCN research relies on simulations in order to test the effects of caching within a provider’s network.

Network topologies are collected and made publicly available by the Rocketfuel project [26,22]. From these, we selected the *Abilene* network since it has plausible topological properties of a future CCN network [22]. From a practical perspective (i.e., to bound the run time of our simulations), Abilene is also suitable due to its small number of routers. Figure 4 illustrates this topology, where the nodes n_0, \dots, n_{10} are routers. We note that it is not our goal to illustrate the advantage of a specific new caching strategy for various network topologies or sizes. Instead, our aim is to provide a flexible, elegant means to tune the behaviour and performance of a given network, and give a proof-of-concept evaluation of our architecture. Carrying out the considered evaluations on multiple networks is beyond the scope of the present feasibility study.

After choosing the Abilene network topology, we connected (virtual) content consumers and producers to the CCN routers. For every simulation run, we connect each of the u consumers uniformly at random to one of the 11 routers. Likewise, all content providers get connected to exactly one router.

Scenarios. As mentioned previously, one of the goals is to adequately react to changes in content popularity scenarios. In our simulations, α is the only system parameter that is varied during the simulations to emulate changes in the users’ content access pattern. To this end, we consider the following two *scenarios*:

- *LHL*: This scenario starts with a low α value, then changes to a high value, then changes back to low.
- *HLH*: Dually, this scenario starts with a high α value, changes then to low and back to high.

The parameter α representing content popularity is central to the caching performance. However, there is no consent in the CCN literature about the exact α . We take the extremal values found in [22,24,27], i.e.,

$$\alpha = 0.4 \text{ (Low) and } \alpha = 2.5 \text{ (High)}.$$

The total time span of each simulation is 1800 seconds, and we switch the value of α after 600 and 1200 seconds. In each of these 600-seconds intervals, each consumer starts downloading a video at a time point selected uniformly at random.

The scenario LHL allows to study how the caching system will react if a small set of contents suddenly becomes very popular. Then, after a certain amount of time, content becomes more equally distributed again. Scenario HLH tests the dual case, where content popularity is more concentrated at the beginning and at the end of the simulation, but less so in the middle.

Caching Strategies (Cache Replacement Policies). As far as basic cache strategies are concerned, we limit the study to Least Frequently Used and Random. Note that other cache replacement policies might provide a higher hit ratio in general, and in particular in the studied setup. However, our goal is neither a study of given caching mechanisms as such, nor the development of a new static caching strategy. Instead, we are interested in evaluating (i) our architecture for flexible configuration based on stream reasoning techniques, and (ii) our hypothesis that intelligent switching between strategies locally depending on the situation may lead to better performance. We review a few observations from [27] on Random and LFU that will also be confirmed in this study.

Static Strategies. The cache replacement policies Random and LFU are used in two ways: as static strategies, and as basic mechanisms to switch between in dynamic strategies.

- *Random.* The Random strategy allows for efficient cache management due to the constant runtime of its replacement operations, leaving computing time for other router tasks. Random usually leads to fewer cache hits than popularity-based replacement policies like LFU or LRU. However, it is a suitable and cheap alternative when content popularity is more equally distributed.
- *Least Frequently Used (LFU).* This strategy offers good cache hit performance when every router caches all forwarded Data packets. However, it needs $\mathcal{O}(\log n)$ comparisons per cache hit, where n is the number of cached items. It is more suitable for a stronger concentration of content popularity but reacts slowly to changes.

Dynamic Strategies. The basic strategies Random and LFU are dynamically employed by the following alternating strategies Admin and Intelligent Caching Agent (ICA).

- *Admin.* The Admin strategy is a hypothetical strategy for testing purposes, where strategy changes are manually configured. In Admin mode, all routers change their replacement policy after 600 or 1200 seconds, respectively, simultaneously with the change of content popularity (as induced by the simulation). That is, each router uses Random replacement in intervals with low α (L phases), and LFU in intervals with high α (H phases).
- *Intelligent Caching Agent (ICA).* In contrast to Admin, ICA is not manually configured and does not enforce the same replacement policy for all routers. ICA switches

the caching mechanism separately for each router based on reasoning on the locally observed data stream. Consider, for instance, a phase with high α , and two routers R_1 and R_2 , such that the more popular content is requested mostly from R_1 . Hence, R_2 observes a more equal distribution and thus has no reason to switch from Random to LFU. Due to this flexibility to react to changing demands, ICA should give a measurable benefit.

Simulation System Parameters. Table 1 lists the parameters used in the evaluation,

Parameter	Values
Simulation duration - t	1800 sec
Number of consumers - u	1000
Number of videos - v	50
Video bit rate - r	1.33 Mbit/s
Video size - V	10 MB (10^3 chunks, duration 60s)
Chunk size - K	10 kB
Cache size - C	[50, 250, 500, 2000, 5000] chunks
Cache percentage - $p = C/(v \cdot V)$	[0.1, 0.5, 1, 4, 10] %

Table 1: Selected simulation parameters

along with their values. We will now explain these parameters in detail.

- *Videos*. The chosen file size V of the distributed videos is 10 MB, which is about the size of an average YouTube video [22]. All $v = 50$ videos in the simulation have a duration of 60 seconds and a constant bit rate of $r = 1.33$ Mbit/s to simplify the simulation setup.
- *Chunk Size - K* . In CCN, the *chunk size* is an important system parameter [22] because it defines both the maximum size of *Data* packets and the size of the cached content pieces (chunks). Any transmitted content is partitioned into chunks of a fixed size. The smaller the chunk size, the more *Interest* packets need to be issued for a single file. Consequently, if the chunk size is too small, the number of requests per *Interest* packet becomes too large. On the other hand, some upper bound on the size of a chunk is necessary. In general, storing entire large files as chunks is not possible (e.g., due to hardware limits), not practical, and invalidates the conceptual approach of CCN [22]. In their seminal paper [18], Jacobson et al. propose packet-size chunks. According to [22], chunk sizes smaller than 10 kB are likely to cause too much overhead.
- *Cache Size - C* . The *cache size* is the maximum size for all caches in a CCN router. It is also a central parameter in the evaluation of the caching mechanism. The cache is limited by its requirement to operate at line speed [19,25]. Furthermore, caches are technically restricted by the memory access latencies of the underlying memory technology. Dynamic Random Access Memory (DRAM) technology is able to deliver the line speed requirement at about 10 GB of storage [19,25] and is used as the upper bound of the cache size in our evaluation. Recent works [19,25] propose the combination of multiple memory technologies with different speeds. Past chunk access patterns are

used to predict future requests for subsequent chunks of the same content [19]. This allows the system to move batches of chunks from the slow SSD storage into the DRAM to avoid the bottleneck of the storage access latency of the SSD [19]. Since our focus is on analyzing the effect of switching caching strategies based on a reaction to the recent content interest distribution, we employ only a single cache size. In order to cope with the overall time of simulation runs, we have to further limit the cache size.

- **Cache Percentage - p .** By the *cache percentage* $p = C/(v \cdot V)$ we understand the relative size of content that can be kept in the *Content Store* of a router. More precisely, p is the ratio of the cache size C and the total size of all of videos. Evidently, p is another central parameter w.r.t. caching performance, and a wide range of values has previously been considered [22]. Clearly, the more content can be cached, the more cache hits will arise for all strategies, and the less difference will emerge between them. However, caching a large proportion of all the entire content catalog is unrealistic. Thus, in terms of evaluating caching strategies for future real-world deployments, it is important to use a realistic (small) caching percentage. Following [22], we target values for p in the range of 10^{-3} to 10^{-1} . For the simulation we first fix the catalog size $v \cdot V$ for which a simulation run can be completed in reasonable time. Then for each run a cache percentage p is assigned by varying the cache size C .

5.2 Performance Metrics

We use the cache hit ratio and the cache hit distance as performance metrics. Both metrics are typically used by the CCN community for evaluating caching mechanisms [31].

Cache Hit Ratio. The *cache hit ratio* is defined as h/req , where h is the number of cache hits and req is the number of requests. A *cache hit* is counted when an *Interest* packet can be satisfied by some router's *Content Store*. In addition to the total hit ratio, we also measure the development of the cache hits over time in intervals of one second. This allows us to get a detailed insight into the system's performance at every time point. The better the cache system works, the higher the hit ratio will be, which will result in a reduction of the network load and the access latencies [31].

Cache Hit Distance. The *cache hit distance* is the average number of *hops* for a *Data* packet, i.e., the number of routers travelled between the router answering a request and the consumer that had issued it. A smaller value is preferable because it will result in a reduction of access latencies.

5.3 Execution

The simulation was executed in two steps. First, we needed to determine a reasonable cache size based on basic caching strategies. After fixing the cache size, we then proceeded with a detailed analysis of different static and dynamic caching strategies.

Step 1: Determining the Cache Size. In the first step, on the selected network topology and for each of the cache percentages ($p = 0.1, 0.5, 1, 4, 10$) we executed 30 times 2 basic strategies (LFU, Random) in 1 Scenario (LHL). The goal of all 300 individual runs was to evaluate how the cache hit ratio behaves in relation to the cache size.




Fig. 5: Hit ratio for basic strategies over various cache sizes

Figure 5 shows the aggregated cache hit ratios of the 30 independent runs each for the LFU and Random caching strategies, respectively. Clearly, the cache size is the dominating parameter w.r.t. the cache hit ratio, i.e., the latter increases with the cache size. Cache sizes far beyond 1% are practically irrelevant due to hardware constraints. Furthermore, switching strategies is then of little interest as LFU anyway dominates in these cases. For cache sizes smaller than 1%, both tested strategies resulted in very low cache hit rates.

The intriguing question is whether switching intelligently between LFU and Random may give a benefit, under realistic conditions where their single static use gives a comparable performance. Thus, the following simulations focused on setups with $p = 1$.

Step 2: Performance Comparison. To evaluate potential performance gains by dynamically switching LFU and Random after fixing $p = 1\%$, we ran 30 tests for every combination of one of the 4 caching strategies (Random, LFU, Admin, ICA) and 2 scenarios (LHL and HLH), i.e., a total of 240 individual runs. In addition to the strategies Random and LFU used above, we considered the Admin strategy, and the Intelligent Caching Agent (ICA) strategy. Recall that the latter two both switch between LFU and Random. The manually configured Admin strategy changes the strategy globally, i.e., for all routers at seconds 600 and 1200. On the other hand, ICA might run a different strategy on different routers, based on the recent content access data which is observed locally.

5.4 Results

We are now going to present the findings of our simulations. First, we will analyze the behaviour of the considered strategies during a single run. Then, we will investigate the effect of switching strategies.




Fig. 6: Cache hits in simulation run for LHL




Fig. 7: Cache Strategies used by ICA for a simulation run

Simulation Run Analysis Reacting to Changing Content Access. Figure 6 shows the development of the cache hits over time during a single simulation run for Scenario LHL. We first note that the number of active users (right y-axis) initially increases rapidly and is then varying only slightly with an average of about 95. Thus, the sudden increase in cache hits between seconds 600 and 700, and the sudden decrease after second 1200, cannot be attributed to a changing number of users. Clearly, these changes are due to the induced switch in content interest patterns, first from the initial L phase (low value α) to the middle H phase (high value α) and then back from H to L.

The cache hit increase after 600 seconds is observed for all caching strategies. Dually, one can see the reverse effect after switching back to a more equal content interest distribution after 1200 seconds. The Random strategy is in general slowly responding to the new situation and shows a steady increase in the middle H phase compared to the rapid increase of hits seen with the other strategies. Note that the Random strategy stores and replaces arbitrary chunks. If requested content becomes less equally distributed, the stored chunks tend to be more often those that are more popular. This explains why cache hits are increasing also under Random and why its reaction is slower.




Fig. 8: Aggregated evaluation results over 30 runs for each caching strategy using $p = 1\%$

LFU reacts well to the change from L to H but still has to deal with the recent history of cache items gathered in the phase with low α . Thus, it does not achieve as high hit rates as the alternating cache strategies Admin and ICA.

It is no surprise that the hypothetical Admin strategy shows a very fast reaction in both situations where the value α changes. Interestingly, the reactive ICA strategy shows about the same performance as Admin. We will now investigate a run of ICA in more detail.

Strategy Alternation by ICA. Figure 7 depicts which of the two basic strategies is used by ICA at which time during the scenario LHL discussed above. Before the first run of the reasoner, all routers (running ICA) start with the LFU strategy and then detect the low value of α . Consequently, they switch to the more suitable Random caching strategy. During the following phase with a high α value, most of the routers switch to LFU after a short delay and stay with this strategy for most of the phase. The fact that not all routers behave equally reflects the random nature of our simulation as well as effects arising from the network topology: not all routers observe the same content access distribution. As we will see below, the flexibility of local decision making gives an advantage over the Admin strategy, which reflects a human user’s decision making for an entire network. The detection of the third phase with a low α value also works well, where all routers switch back to Random caching, i.e., the superior strategy for more equal distribution of interest.

Performance Comparison Figure 8 presents an overview of our performance comparisons, indicating the performance of caching strategies LFU, Admin and ICA in relation to Random, which is used as baseline (100%). Figures 8a and 8b depict the obtained cache hit ratios, and Figures 8c and 8d show the results for cache hit distances. All box plots visualize the aggregated results over 30 individual runs with the respective caching strategy, i.e., LHL for Figures 8a and 8c and HLH for Figures 8b and 8d. We are now going to analyze the obtained results.

Cache Hit Ratios. Figure 8a compares the cache hit ratios of all strategies in scenario LHL. We see that LFU has a higher variance than Random and is slightly worse on average. Both dynamic strategies, Admin and ICA, which switch between LFU and Random, outperform the static approaches.

Figure 8b depicts the converse scenario HLH. In the two H phases, which comprise two thirds of the overall runtime, LFU works well and thus the ratios are closer to each other than in LHL. Still, ICA is performing better than the other strategies, even compared to Admin. Here, the benefit of separate strategies for routers arises: in contrast to Admin, which switches all routers at a phase change, each router individually decides in ICA by reasoning about the locally observed data stream.

Both plots (8a and 8b) demonstrate that intelligent alternation between caching strategies results in better performance. Moreover, ICA shows significantly less variation than LFU in both scenarios.

Cache Hit Distances. Similarly as for cache hit ratios, Figures 8c and 8d show the aggregated cache hit distances for the considered caching strategies for scenarios LHL and HLH, respectively.

Also according to this metric, Admin and ICA deliver better performance than the static approaches. Figure 8c again shows a clear difference between static and alternating strategies, in terms of mean values and of variance. Figure 8d confirms the better performance of LFU compared to Random for the Scenario HLH. As for cache hit ratio, LFU is also close to the Admin strategy in terms of cache hit distance. Finally, as above, ICA is even better than Admin due to its flexibility.

In summary, dynamic switching is advantageous in both settings. ICA is as least as good as Admin in LHL, and proves to be the best strategy for HLH. Notably, both dynamic strategies lead to a decreased cache hit distance relative to the Random strategy.

6 Conclusion

We presented a comprehensive feasibility study how reasoning techniques can be used for adaptive configuration tasks that depend on streaming data. More specifically, we provided an architecture for the simulation of potential caching strategies in future Content-Centric Networks. Our empirical evaluations indicate that dynamic switching of caching strategies in reaction to changing user behavior may give significant savings due to performance gains.

We focused on a principled approach of automated decision making by means of high-level reasoning on stream data and provided a purely declarative control unit. To obtain a program from our formal LARS models, we implemented plug-ins for DLVHEX. The resulting encoding resembles the mathematical counterpart, i.e., a fragment of LARS. Notably, full-scale industrial solutions will involve much more complex decision rules and processes, and fast empirical assessment will be crucial. Thus, it would be very valuable to have tools such that the formal modeling directly gives us an executable specification, as in the case of Answer Set Programming.

These observations clearly motivate the advancement of stream reasoning research, especially on the practical side. In particular, stream processing engines are in need that have an expressive power similar to LARS. While a lot of resources are currently being invested into efficient, distributed, low-level processing of so-called Big Data, declarative methods to obtain traceable insights from streams need more attention.

References

1. Arasu, A., Babu, S., Widom, J.: The CQL continuous query language: semantic foundations and query execution. *VLDB J.* 15(2), 121–142 (2006)
2. Babu, S., Widom, J.: Continuous queries over data streams. *SIGMOD Record* 3(30), 109–120 (2001)
3. Beck, H., Dao-Tran, M., Eiter, T., Fink, M.: LARS: A Logic-based Framework for Analyzing Reasoning over Streams. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, January 25–29, 2015, Austin, Texas, USA (2015)
4. Bernardini, C., Silverston, T., Festor, O.: MPC: popularity-based caching strategy for content centric networks. In: *Proceedings of IEEE International Conference on Communications, ICC 2013, Budapest, Hungary, June 9–13, 2013*. pp. 3619–3623. IEEE (2013), <http://dx.doi.org/10.1109/ICC.2013.6655114>
5. Brewka, G., Eiter, T., Truszczyński, M.: Answer set programming at a glance. *Communications of the ACM* 54(12), 92–103 (2011)
6. Cha, M., Kwak, H., Rodriguez, P., Ahn, Y., Moon, S.B.: Analyzing the video popularity characteristics of large-scale user generated content systems. *IEEE/ACM Trans. Netw.* 17(5), 1357–1370 (2009), <http://doi.acm.org/10.1145/1665838.1665839>
7. Cho, K., Lee, M., Park, K., Kwon, T.T., Choi, Y., Pack, S.: WAVE: popularity-based and collaborative in-network caching for content-oriented networks. In: *2012 Proceedings IEEE INFOCOM Workshops, Orlando, FL, USA, March 25–30, 2012*. pp. 316–321. IEEE (2012), <http://dx.doi.org/10.1109/INFCOMW.2012.6193512>
8. Cisco: Cisco Visual Networking Index: Forecast and Methodology, 2014–2019. White Paper pp. 2014–2019 (2016), <http://www.cisco.com/c/en/us/solutions/collateral/>

- service-provider/ip-ngn-ip-next-generation-network/
white({})paper({})c11-481360.pdf
9. Della Valle, E., Ceri, S., van Harmelen, F., Fensel, D.: It's a Streaming World! Reasoning upon Rapidly Changing Information. *IEEE Intelligent Systems* 24, 83–89 (2009)
 10. Do, T.M., Loke, S.W., Liu, F.: Answer Set Programming for Stream Reasoning. In: *Adv. Artif. Intell.* pp. 104–109 (2011)
 11. Eiter, T., Fink, M., Krennwallner, T., Redl, C.: Domain Expansion for ASP-Programs with External Sources. *Artif. Intell.* 233, 84–121 (2014), <http://dx.doi.org/10.1016/j.artint.2016.01.003>
 12. Eiter, T., Mehuljic, M., Redl, C., Schüller, P.: User guide: dlhex 2.x. Tech. Rep. INF-SYS RR-1843-15-05, Vienna University of Technology, Institute for Information Systems (September 2015)
 13. Faber, W., Leone, N., Pfeifer, G.: Recursive aggregates in disjunctive logic programs: Semantics and complexity. In: *JELIA*. pp. 200–212 (2004), http://dx.doi.org/10.1007/978-3-540-30227-8_19
 14. Gebser, M., Grote, T., Kaminski, R., Obermeier, P., Sabuncu, O., Schaub, T.: Stream Reasoning with Answer Set Programming. Preliminary Report. In: *KR*. pp. 613–617 (2012)
 15. Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., Thiele, S.: Engineering an incremental ASP solver. In: *ICLP*. pp. 190–205 (2008)
 16. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. *New Gener. Comput.* 9(3-4), 365–386 (1991), <http://link.springer.com/article/10.1007/BF03037169>
 17. Handley, M.: Why the internet only just works. *BT Technology Journal* 24(3), 119–129 (Jul 2006), <http://dx.doi.org/10.1007/s10550-006-0084-z>
 18. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.: Networking named content. In: Liebeherr, J., Ventre, G., Biersack, E.W., Keshav, S. (eds.) *Proceedings of the 2009 ACM Conference on Emerging Networking Experiments and Technology, CoNEXT 2009, Rome, Italy, December 1-4, 2009*. pp. 1–12. *ACM* (2009), <http://doi.acm.org/10.1145/1658939.1658941>
 19. Mansilha, R.B., Saino, L., Barcellos, M.P., Gallo, M., Leonardi, E., Perino, D., Rossi, D.: Hierarchical content stores in high-speed ICN routers: Emulation and prototype implementation. In: Solis, I., Carzaniga, A., Ramakrishnan, K.K. (eds.) *Proceedings of the 2nd International Conference on Information-Centric Networking, ICN '15, San Francisco, California, USA, September 30 - October 2, 2015*. pp. 59–68. *ACM* (2015), <http://doi.acm.org/10.1145/2810156.2810159>
 20. Mastorakis, S., Afanasyev, A., Moiseenko, I., Zhang, L.: ndnSIM 2.0: A new version of the NDN simulator for NS-3. Technical Report NDN-0028, NDN (January 2015)
 21. Mileo, A., Abdelrahman, A., Policarpio, S., Hauswirth, M.: Streamrule: A nonmonotonic stream reasoning system for the semantic web. In: Faber, W., Lembo, D. (eds.) *Web Reasoning and Rule Systems - 7th International Conference, RR 2013, Mannheim, Germany, July 27-29, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7994*, pp. 247–252. *Springer* (2013), <http://dx.doi.org/10.1007/978-3-642-39666-3>
 22. Rossi, D., Rossini, G.: Caching performance of content centric networks under multi-path routing (and more). Relatório técnico, Telecom ParisTech (2011)
 23. Rossi, D., Rossini, G.: On sizing CCN content stores by exploiting topological information. In: *2012 Proceedings IEEE INFOCOM Workshops, Orlando, FL, USA, March 25-30, 2012*. pp. 280–285 (2012), <http://dx.doi.org/10.1109/INFCOMW.2012.6193506>
 24. Rossini, G., Rossi, D.: A dive into the caching performance of content centric networking. In: *17th IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks, CAMAD 2012, Barcelona, Spain, September 17-19, 2012*. pp. 105–109. *IEEE* (2012), <http://dx.doi.org/10.1109/CAMAD.2012.6335307>

25. Rossini, G., Rossi, D., Garetto, M., Leonardi, E.: Multi-terabyte and multi-gbps information centric routers. In: 2014 IEEE Conference on Computer Communications, INFOCOM 2014, Toronto, Canada, April 27 - May 2, 2014. pp. 181–189. IEEE (2014), <http://dx.doi.org/10.1109/INFOCOM.2014.6847938>
26. Spring, N.T., Mahajan, R., Wetherall, D., Anderson, T.E.: Measuring ISP topologies with rocketfuel. *IEEE/ACM Trans. Netw.* 12(1), 2–16 (2004), <http://dx.doi.org/10.1109/TNET.2003.822655>
27. Tarnoi, S., Suksomboon, K., Kumwilaisak, W., Ji, Y.: Performance of probabilistic caching and cache replacement policies for content-centric networks. In: IEEE 39th Conference on Local Computer Networks, LCN 2014, Edmonton, AB, Canada, 8-11 September, 2014. pp. 99–106. IEEE Computer Society (2014), <http://dx.doi.org/10.1109/LCN.2014.6925761>
28. Xylomenos, G., Ververidis, C.N., Siris, V.A., Fotiou, N., Tsilopoulos, C., Vasilakos, X., Katsaros, K.V., Polyzos, G.C.: A survey of information-centric networking research. *IEEE Communications Surveys and Tutorials* 16(2), 1024–1049 (2014), <http://dx.doi.org/10.1109/SURV.2013.070813.00063>
29. Yu, H., Zheng, D., Zhao, B.Y., Zheng, W.: Understanding user behavior in large-scale video-on-demand systems. In: Berbers, Y., Zwaenepoel, W. (eds.) *Proceedings of the 2006 EuroSys Conference*, Leuven, Belgium, April 18-21, 2006. pp. 333–344. ACM (2006), <http://doi.acm.org/10.1145/1217935.1217968>
30. Zaniolo, C.: Logical foundations of continuous query languages for data streams. In: *Data-log*. pp. 177–189 (2012)
31. Zhang, M., Luo, H., Zhang, H.: A Survey of Caching Mechanisms in Information-Centric Networking. *IEEE Communications Surveys and Tutorials* 17(3), 1473–1499 (2015), <http://dx.doi.org/10.1109/COMST.2015.2420097>